arXiv:1510.06587v2 [cs.MA] 30 Aug 2016

Approximating Strategic Abilities under Imperfect Information: a Naive Approach

Wojciech Jamroga, Michal Knapik, and Damian Kurpiewski

Institute of Computer Science,
Polish Academy of Sciences, Poland
{w.jamroga,michal.knapik,damian.kurpiewski}@ipipan.waw.pl

Abstract. Alternating-time temporal logic (**ATL**) allows to specify requirements on abilities that different agents should (or should not) possess in a multi-agent system. However, model checking **ATL** specifications in realistic systems is computationally hard. In particular, if the agents have imperfect information about the global state of the system, the complexity ranges from Δ_2^{P} to undecidable, depending on the syntactic and semantic details. The problem is also hard in practice, as evidenced by several recent attempts to tackle it. On the other hand, model checking of alternating epistemic mu-calculus can have a distinctly lower computational complexity. In this work, we look at the idea of approximating the former problem by the verification of its "naive" translations to the latter. In other words, we look at what happens when one uses the (incorrect) fixpoint algorithm to verify formulae of **ATL** with imperfect information.

1 Introduction

There is a growing number of works that study syntactic and semantic variants of *strategic logics*, in particular the alternating-time temporal logic **ATL**. Conceptually, the most interesting strand builds upon reasoning about temporal patterns and outcomes of strategic play, limited by information available to the agents. The contributions are mainly theoretical, and include results concerning the conceptual soundness of a given semantics of ability [17,9,1,13], meta-logical properties [4], and the complexity of model checking [17,11,10]. However, there is very little research on the actual *use* of the logics, in particular on practical algorithms for reasoning and/or verification.

This is somewhat easy to understand, since model checking of **ATL** variants with imperfect information has been proved Δ_2^{P}- to PSPACE-complete for agents playing positional (a.k.a. memoryless) strategies [17,11] and undecidable for agents with perfect recall of the past [7]. Moreover, the imperfect information semantics of **ATL** does not admit fixpoint equivalences [4], which makes incremental synthesis of strategies impossible, or at least difficult to achieve. Some practical attempts at tackling the problem started to emerge only recently [15,5,8]. Up until now, experimental results confirm that the initial intuition was right: model checking strategic modalities for imperfect information is

hard, and dealing with it requires innovative algorithms and verification techniques.

One idea that has not been properly explored is that of alternating-time epistemic mu-calculus (**AE**μ**C**) [3]. Since fixpoint equivalences do not hold under imperfect information, it follows that standard fixpoint translations of **ATL** modalities lead to a different interpretation of strategic ability. In fact, it can be argued that they capture existence of *recomputable* winning strategies. However, what especially interests us in the context of model checking is that they can make model checking computationally cheaper. Verification of **AE**μ**C** is in general between NP and Δ_2^{P}, but the scope of backtracking is much smaller than for **ATL** with imperfect information, as it includes only the actions starting from a given indistinguishability class rather than all the actions in the model. Moreover, for coalitions of up to 2 agents the model checking problem is in P [3].

The question that we ask in this paper is: Is **AE**μ**C** an attractive alternative for verification of strategic abilities under imperfect information? To this end, we will look at the naive **AE**μ**C** approximations of formulae of $\mathbf{ATL}_{\mathrm{ir}}$ (i.e., **ATL** with imperfect information and imperfect recall), and investigate:

1. Whether model checking of the **AE**μ**C** approximations performs significantly faster than for the original $\mathbf{ATL}_{\mathrm{ir}}$ formulae;
2. Whether the **AE**μ**C** counterparts are indeed semantic approximations of the $\mathbf{ATL}_{\mathrm{ir}}$ specifications, or they encapsulate a completely different concept of ability. As fixpoint equivalences are not valid for $\mathbf{ATL}_{\mathrm{ir}}$, we know that the naive fixpoint translation is in general incorrect. However, one can possibly ask: how often?

We take on an empirical approach. More precisely, we consider two classes of benchmark models and formulae, one based on the Tian Ji scenario [14,5] and the other being the Castles benchmark from [15]. Then, for a formula φ, we compare the output and performance of the $\mathbf{ATL}_{\mathrm{ir}}$ model checking of φ with the **AE**μ**C** model checking of $aemc(\varphi)$, i.e., with the straightforward (and generally incorrect) fixpoints approximation of φ. The work reported here is very preliminary, but it already allows to draw some conclusions, and decide on the most promising lines for future research.

2 What Agents Can Achieve under Imperfect Information

In this section we provide a brief overview of the relevant variants of **ATL**, and the corresponding complexity results for model checking. We refer the interested reader to [2,17,3] for details.

2.1 Models

The semantics for **ATL** is defined over a variant of transition systems where transitions are labeled with combinations of actions, one per agent. An *imperfect information concurrent game structure* (ICGS) [2,17] is given by $M =$

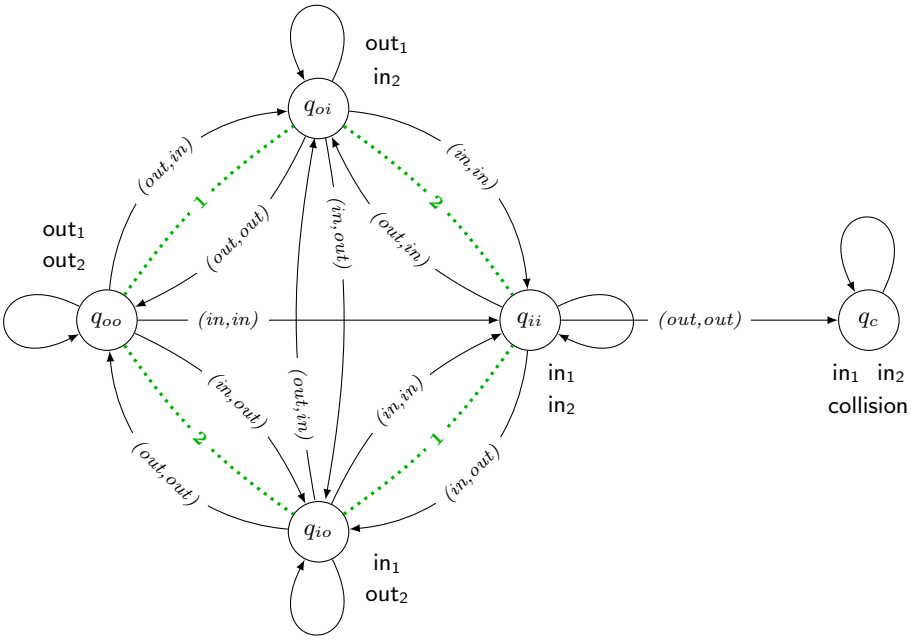

Fig. 1. Autonomous vehicles at the intersection: model M_1

$\langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, o, \{\sim_a | \ a \in \mathbb{A}\mathrm{gt}\}\rangle$ which includes a nonempty finite set of all agents $\mathbb{A}\mathrm{gt} = \{1, \ldots, k\}$, a nonempty set of states St, a set of atomic propositions Π and their valuation $\pi : Prop \to 2^{St}$, and a nonempty finite set of (atomic) actions Act. Function $d : \mathbb{A}\mathrm{gt} \times St \to 2^{Act}$ defines nonempty sets of actions available to agents at each state, and o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \ldots, \alpha_k)$ to state q and a tuple of actions $\langle \alpha_1, \ldots, \alpha_k \rangle$ for $\alpha_i \in d(i, q)$ and $1 \leq i \leq k$, that can be executed by $\mathbb{A}\mathrm{gt}$ in q. We write $d_a(q)$ instead of $d(a, q)$. Each $\sim_a \subseteq St \times St$ is an equivalence relation satisfying $d_a(q) = d_a(q')$ for $q \sim_a q'$. Note that perfect information can be modeled by assuming each $\sim_a$ to be the minimal reflexive relation.

Example 1 (Intersection with limited visibility). Consider an intersection with k autonomous vehicles around it. Each vehicle is modeled as a separate agent, whose local state is characterized by either the proposition $\mathsf{out_i}$ (when the vehicle is outside the intersection) or $\mathsf{in_i}$ (when the vehicle is inside it). The available actions are: *in* ("drive in" or "stay in", depending on the current state) and *out* ("drive out" or "stay out"). Transitions update the state accordingly, except for one case: when both agents are in and decide to leave at the same time, a collision occurs (collision). Furthermore, let us assume that no agent sees the location of the other vehicle.

Figure 1 presents a pointed ICGS modeling the scenario for $k = 2$. The combinations of actions that are not displayed in the graph do not change the state of the system. The indistinguishability relations are depicted by dotted lines. □

A *strategy* of agent a is a conditional plan that specifies what a is going to do in each situation. Here, we only refer to *memoryless uniform strategies* (ir *strategies* in short), defined as functions $s_a : St \to Act$ such that $s_a(q) \in d_a(q)$ for all q, and $q \sim_a q'$ implies $s_a(q) = s_a(q')$. A *collective strategy* s_A is a tuple of ir strategies, one per agent from A. A *path* $\lambda = q_0 q_1 q_2 \dots$ is an infinite sequence of states such that there is a transition between each q_i, q_{i+1}. We use $\lambda[i]$ to denote the ith position on path λ (starting from $i = 0$) and $\lambda[i, \infty]$ to denote the subpath of λ starting from i. Function $out_M(q, s_A)$ returns the set of all paths that can result from the execution of strategy s_A from state q in model M, defined formally as follows:

$out(q, s_A) = \{\lambda = q_0, q_1, q_2 \dots \mid q_0 = q$ and for each $i = 0, 1, \dots$ there exists $\langle \alpha^i_{a_1}, \dots, \alpha^i_{a_k} \rangle$ such that $\alpha^i_a \in d_a(q_i)$ for every $a \in \mathbb{A}\mathrm{gt}$, and $\alpha^i_a = s_A[a](q_i)$ for every $a \in A$, and $q_{i+1} = o(q_i, \alpha^i_{a_1}, \dots, \alpha^i_{a_k})\}$.

Moreover, we define $out^{ir}_M(q, s_A) = \bigcup_{a \in A} \bigcup_{q \sim_a q'} out_M(q', s_A)$. We will omit the subscript M if it is clear from the context.

2.2 Alternating Time Temporal Logic

Let $\mathbb{A}\mathrm{gt}$ be the set of agents and *Prop* be the set of atomic propositions. The language of **ATL** is given by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle \mathsf{X}\varphi \mid \langle\langle A \rangle\rangle \mathsf{G}\varphi \mid \langle\langle A \rangle\rangle \varphi \,\mathsf{U}\, \varphi$$

where $A \subseteq \mathbb{A}\mathrm{gt}$ and $\mathsf{p} \in Prop$. Additionally, we define "sometime in the future" as $\mathsf{F}\varphi \equiv \top \,\mathsf{U}\, \varphi$. The semantics of $\mathbf{ATL}_{\mathrm{ir}}$ is defined by the following clauses:

- $M, q \models p$ iff $q \in \pi(p)$ (where $p \in \Pi$);
- $M, q \models \neg\varphi$ iff $M, q \not\models \varphi$;
- $M, q \models \varphi \wedge \psi$ iff $M, q \models \varphi$ and $M, q \models \psi$;
- $M, q \models \langle\langle A \rangle\rangle \mathsf{X}\varphi$ iff there is a collective strategy s_A such that, for each $\lambda \in out^{ir}(q, s_A)$, $M, \lambda[1] \models \varphi$;
- $M, q \models \langle\langle A \rangle\rangle \mathsf{G}\varphi$ iff there is a collective strategy s_A such that, for each $\lambda \in out^{ir}(q, s_A)$ and every $i \geq 0$, $M, \lambda[i] \models \varphi$;
- $M, q \models \langle\langle A \rangle\rangle \varphi \,\mathsf{U}\, \psi$ iff there is s_A such that, for each $\lambda \in out^{ir}(q, s_A)$, there is $i \geq 0$ for which $M, \lambda[i] \models \psi$, and $M, \lambda[j] \models \varphi$ for each $0 \leq j < i$.

Informally, $M, q \models \langle\langle A \rangle\rangle \gamma$ iff there exists a strategy for A such that γ holds on all the paths that the agents in A consider as possible executions of the strategy.

Example 2 (Intersection with limited visibility, ctd.). Take model M_1 from Example 1. Now, we have e.g. that $M_1, q_{oo} \models \langle\langle 1 \rangle\rangle \mathsf{G}\neg\mathsf{collision}$ (it suffices that agent

1 executes action "*out*" regardless of anything). On the other hand, $M_1, q_{oo} \models \neg\langle\langle 1,2\rangle\rangle \mathsf{Fcollision}$ (the agents do not know how to make sure that a collision will happen, even if they want to). We leave it up to the interested reader to check the latter. □

2.3 Verification of Strategic Abilities

The model checking problem asks, given a model M, a state q in it, and a logical formula φ, whether φ holds in M, q. **ATL** verification is known to be tractable for perfect information models, but intractable for imperfect information.

Proposition 1 ([17,12]). *Model checking of* $\mathbf{ATL}_{\mathrm{ir}}$ *is* Δ_2^{P}*-complete in the number of states and transitions in the model, and the length of the formula.*

2.4 Alternating Epistemic Mu-Calculus

Alternating epistemic μ-calculus ($\mathbf{AE}\mu\mathbf{C}$) replaces the temporal-strategic operators $\langle\langle A\rangle\rangle\mathsf{G}, \langle\langle A\rangle\rangle\,\mathsf{U}$ with the *least fixpoint operator* μ [3]:

$$\varphi ::= p \mid Z \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A\rangle\rangle\mathsf{X}\varphi \mid \mu Z.\varphi$$

where $Z \in \mathcal{V}ar$ is a second order variable ranging over sets of states. The *greatest fixed point operator* ν can be defined as dual to μ. We consider only the *alternation-free* fragment of $\mathbf{AE}\mu\mathbf{C}$, cf. [2,3] for the exact definition.

A *valuation of* $\mathcal{V}ar$ is a mapping $\mathcal{V} : \mathcal{V}ar \to 2^{St}$. Given a set $Q \subseteq St$ of states we define $\mathcal{V}[Z := Q]$ as the update of $\mathcal{V}$ that assigns Q to Z. The semantics of $\mathbf{AE}\mu\mathbf{C}$ is given by the denotation function $[\![\varphi]\!]^M_{\mathcal{V}}$ which defines the set of states in M that satisfy φ in the following way: $[\![\mathsf{p}]\!]^M_{\mathcal{V}} = \pi(\mathsf{p})$, $[\![Z]\!]^M_{\mathcal{V}} = \mathcal{V}(Z)$, $[\![\neg\varphi]\!]^M_{\mathcal{V}} = St \setminus [\![\varphi]\!]^M_{\mathcal{V}}$, $[\![\varphi \wedge \psi]\!]^M_{\mathcal{V}} = [\![\varphi]\!]^M_{\mathcal{V}} \cap [\![\psi]\!]^M_{\mathcal{V}}$, $[\![\langle\langle A\rangle\rangle\mathsf{X}\varphi]\!]^M_{\mathcal{V}} = \{q \mid \exists \alpha_A \in d_A(q) \forall \alpha_{\mathbb{A}\mathrm{gt}\setminus A} \in d_{\mathbb{A}\mathrm{gt}\setminus A} : o(q, (\alpha_A, \alpha_{\mathbb{A}\mathrm{gt}\setminus A})) \in [\![\varphi]\!]^M_{\mathcal{V}}\}$, and $[\![\mu Z.\varphi]\!]^M_{\mathcal{V}} = \bigcap\{Q \subseteq St \mid [\![\varphi]\!]^M_{\mathcal{V}[Z:=Q]} \subseteq Q\}$. Moreover, $M, q \models \varphi$ iff $q \in [\![\varphi]\!]^M_{\mathcal{V}}$ for all valuations $\mathcal{V}$.

Proposition 2 ([3]). *Model checking of* $\mathbf{AE}\mu\mathbf{C}$ *is* P*-complete if all the coalitions in* φ *consist of at most 2 agents. For abilities of coalitions with* 3 *or more agents, it is between* NP *and* Δ_2^{P} *in the size of the largest abstraction class of relations* $\sim_1, \ldots, \sim_k$.

Thus, verification of $\mathbf{AE}\mu\mathbf{C}$ is potentially more attractive than $\mathbf{ATL}_{\mathrm{ir}}$. A natural idea is to use the naive translation of $\mathbf{ATL}_{\mathrm{ir}}$ to $\mathbf{AE}\mu\mathbf{C}$, defined as follows:

$$\begin{aligned}
aemc(p) &= p\\
aemc(\neg\varphi) &= \neg aemc(\varphi)\\
aemc(\varphi \wedge \psi) &= aemc(\varphi) \wedge aemc(\psi)\\
aemc(\langle\langle A\rangle\rangle\mathsf{X}\varphi) &= \langle\langle A\rangle\rangle\mathsf{X}aemc(\varphi)\\
aemc(\langle\langle A\rangle\rangle\mathsf{G}\varphi) &= \nu Z.(aemc(\varphi) \wedge \langle\langle A\rangle\rangle\mathsf{X}Z)\\
aemc(\langle\langle A\rangle\rangle\mathsf{F}\varphi) &= \mu Z.(aemc(\varphi) \vee \langle\langle A\rangle\rangle\mathsf{X}Z)\\
aemc(\langle\langle A\rangle\rangle\varphi\,\mathsf{U}\,\psi) &= \mu Z.(aemc(\psi) \vee aemc(\varphi) \wedge \langle\langle A\rangle\rangle\mathsf{X}Z).
\end{aligned}$$

We will look at how it works in the subsequent sections.

Note that, in a way, $aemc(\langle\langle A\rangle\rangle \mathsf{F}\varphi)$ expresses a *persistent* ability to achieve φ. Likewise, $aemc(\langle\langle A\rangle\rangle \mathsf{G}\varphi$ expresses a persistent ability to maintain φ. This is because $aemc(\langle\langle A\rangle\rangle \mathsf{F}\varphi)$ produces a strategy for A such that $aemc(\langle\langle A\rangle\rangle \mathsf{F}\varphi)$ will also hold for every state reachable by the strategy. For $aemc(\langle\langle A\rangle\rangle \mathsf{G}\varphi)$ the situation is analogous.

3 Fixpoint Approximation of Strategic Ability: Performance

In this section, we empirically compare the performance of model checking $\mathbf{ATL}_{\mathrm{ir}}$ specifications vs. their naive $\mathbf{AE}\mu\mathbf{C}$ approximations.

The experimental results on the $\mathbf{AE}\mu\mathbf{C}$ side have been obtained by running a straightforward implementation of the fixpoint model checking algorithm, implemented in Python 3. The tests have been conducted on a MacBook with an Intel Core i5 CPU with dynamic clock speed of 1.4 GHz, 4 GB of RAM (one module DDR3, 1600 MHz buz clock), and OS X 10.10.5 Yosemite.

The performance of $\mathbf{ATL}_{\mathrm{ir}}$ model checking for the Castles model is cited after [16], and was obtained on a notebook with an Intel Core i7-3630QM CPU with dynamic clock speed of 2.4 GHz up to 3.4 GHz, and 8 GB of RAM (two modules DDR3 PC3-12800, 800 MHz bus clock, effective data rate 1600 MT/s, in dual-channel configuration). Two model checkers were used: SMC and MCMAS. The experiments with SMC were conducted on Windows 7 OS, the experiments with MCMAS on Linux Ubuntu 12.04.2. Thus, the $\mathbf{ATL}_{\mathrm{ir}}$ model checking for Castles was performed on a significantly better computing equipment than the $\mathbf{AE}\mu\mathbf{C}$ verification. We also note that SMC uses several reduction techniques to restrict the search space, and MCMAS operates on compact symbolic representations of models, based on BDD's. In contrast, our model checking of $\mathbf{AE}\mu\mathbf{C}$ was done with a straightforward implementation of the standard explicit state algorithm with no optimizations at all.

The performance of $\mathbf{ATL}_{\mathrm{ir}}$ model checking for the TianJi model is cited after [5], and was obtained with an experimental model checker implemented with PyNuSMV, a Python framework for prototyping and experimenting with BDD-based model-checking algorithms based on NuSMV [6].[1] Besides compact symbolic representations of states and transitions, the model checker features multiple optimization techniques. The authors do not describe the computing configuration that was used for their experiments.

The timeout in all cases is defined as 120 minutes.

3.1 Benchmark 1: Castles

The *Castles* model have been proposed in [15]. The model consists of one agent called *Environment* that keeps track of the health points of three castles, plus

[1] It should be mentioned that the results for TianJi in [5] were obtained for a slightly different semantics of $\mathbf{ATL}_{\mathrm{ir}}$, employing additional fairness constraints.

	AEμC (ψ_1')			**ATL**$_{\text{ir}}$/SMC(ψ_1)	**ATL**$_{\text{ir}}$/MCMAS(ψ_1)
Configuration	**time**	**#sat**	**#iter**	**time**	**time**
4 (1,1,1)	0.011	128	1	timeout	72
5 (1,1,2)	0.024	256	1	timeout	timeout
6 (2,1,2)	0.386	512	1	timeout	timeout
7 (2,2,2)	9.231	1024	1	timeout	timeout
8 (3,2,2)	4352.891	5504	2	timeout	timeout

Fig. 2. Model checking performance for Castles, formula ψ_1' vs. ψ_1

	AEμC (ψ_2')			**ATL**$_{\text{ir}}$/SMC(ψ_2)	**ATL**$_{\text{ir}}$/MCMAS(ψ_2)
Configuration	**time**	**#sat**	**#iter**	**time**	**time**
4 (1,1,1)	0.004	8	1	timeout	78
5 (1,1,2)	0.015	16	1	timeout	error
6 (2,1,2)	0.050	32	1	?	?
7 (2,2,2)	0.225	64	1	?	?
8 (3,2,2)	1.202	128	1	?	?

Fig. 3. Model checking performance for Castles, formula ψ_2' vs. ψ_2. "?" indicates lack of data about the performance of the given model checker on the given instance.

a number of agents called *Workers* each of whom works for the benefit of a castle. Health points (HP, ranging from 0 to 3) represent the current condition of the castle; 0 HP means that the castle is *defeated*. Workers can execute the following actions: attack a castle they do not work for, defend the castle they do work for, or do nothing. Doing nothing is the only available action to a Worker of a defeated castle. No agent can defend its castle twice in a row, it must wait one step before being able to defend again. A castle gets damaged if the number of attackers is greater than the number of defenders, and the damage is equal to the difference. In the initial state, all the castles have 3 HP and every Worker can engage in defending its castle. The indistinguishability relations for Workers are defined as follows. Every Worker knows if it can currently engage in defending its castle, and can observe for each castle if it is defeated or not. The model is parameterized by the number of agents and the allocation of Workers. For example, an instance with 1 worker assigned to the first castle, 3 workers assigned to the second and 4 to the third castle will be denoted by *9 (1,3,4)*.

Formulae We considered the following formulas for Castles:

$$\psi_1 \equiv \langle\langle c12 \rangle\rangle \mathsf{F} \mathit{castle3defeated}$$
$$\psi_2 \equiv \langle\langle w12 \rangle\rangle \mathsf{F} \mathit{allDefeated}$$

The first formula says that the workers working for castles 1 and 2 have a collective strategy to defeat castle 3, no matter what other agents do. Similarly, the second formula says that workers number 1 and 2 have a collective strategy to

	AEμC (ϕ_1')			**ATL$_{\text{ir}}$** (ϕ_1)
Horses	**time**	**#sat**	**#iter**	**time**
3	0.0002	11	2	2.603
4	0.001	18	2	8.205
5	0.014	153	3	30.885
6	0.024	300	3	99.931
7	0.753	2258	4	586.126
8	6.204	4900	4	?

Fig. 4. Model checking performance for TianJi, formula ϕ_1' vs. ϕ_1

	AEμC (ϕ_2')			**ATL$_{\text{ir}}$** (ϕ_2)
Horses	**time**	**#sat**	**#iter**	**time**
3	0.0002	0	2	≈ 2.5
4	0.001	0	2	≈ 10
5	0.003	0	2	≈ 650
6	0.004	0	2	?
7	0.016	0	2	?
8	0.075	0	2	?

Fig. 5. Model checking performance for TianJi, formula ϕ_2' vs. ϕ_2

ensure the defeat of all castles. After the naive translation to **AEμC** we obtain:

$$\psi_1' \equiv \mu Z.(\mathit{castle3defeated} \vee \langle\langle c12 \rangle\rangle \mathsf{X} Z)$$
$$\psi_2' \equiv \mu Z.(\mathit{allDefeated} \vee \langle\langle w12 \rangle\rangle \mathsf{X} Z)$$

Experimental Results The results for Castles are presented in Figures 2 and 3. The tables present results for a sequence of models of various sizes. The column headers are interpreted as follows:

- time: model checking time (in seconds),
- #sat: number of states in which the formula is satisfied,
- #iter: number of iterations until reaching fixpoint. Note: we only have data about #sat and #iter for the experiments with **AEμC** model checking.

3.2 Benchmark 2: TianJi

The second series of experiments has been conducted for the TianJi variant from [5]. The model consists of two agents: Tian Ji and the king. Each agent has n horses numbered $1, \ldots, n$. In the game, Tian Ji and the king send their horses one by one against each other. Horse i of Tian Ji wins the race with king's horse j iff $i > j$. At each stage, the agents know the current score and their own remaining horses, but not those of the opponent. Moreover, the decisions at each round are made simultaneously, so one does not know which horse is currently sent by the other player. The player whose horses won most races wins the game.

Configuration	$\mathbf{AE}\mu\mathbf{C}$ (ψ_1')	$\mathbf{ATL}_{\mathrm{ir}}$ (ψ_1)
4 (1,1,1)	false	true
5 (1,1,2)	false	true
6 (2,1,2)	false	true
6 (3,1,1)	true	true
7 (2,2,2)	false	true
8 (3,2,2)	true	true

Fig. 6. Model checking output for Castles, formula ψ_1' vs. ψ_1

We considered the following formulas for TianJi:

$$\phi_1 := \langle\!\langle TJ \rangle\!\rangle \mathsf{F}\, TJWins$$
$$\phi_2 := \langle\!\langle TJ \rangle\!\rangle \mathsf{G} \langle\!\langle TJ \rangle\!\rangle \mathsf{X}\, TJWonLess2$$

TJWins holds when the game is done and TianJi has won more races than the king. Similarly, *TJWonLess2* is satisfied when TianJi has won at most 1 race up to the current point. After the naive translation we obtain:

$$\phi_1' := \mu Z.(TJWins \vee \langle\!\langle TJ \rangle\!\rangle \mathsf{X} Z)$$
$$\phi_2' := \nu Z.(\langle\!\langle TJ \rangle\!\rangle \mathsf{X}\, TJWonLess2 \wedge \langle\!\langle TJ \rangle\!\rangle \mathsf{X} Z)$$

The results of experiments are presented in Figures 4 and 5.

3.3 Discussion

The experimental results show that, for the instances of model checking that we have tested, verification of strategic abilities in alternating epistemic μ-calculus offers a dramatic speedup over model checking "standard" **ATL** with imperfect information. So, from the computational point of view, $\mathbf{AE}\mu\mathbf{C}$ is definitely more attractive than $\mathbf{ATL}_{\mathrm{ir}}$. The speedup occurs despite the fact that the $\mathbf{AE}\mu\mathbf{C}$ model checking was done by a straightforward implementation of the standard explicit state algorithm, while the results for $\mathbf{ATL}_{\mathrm{ir}}$ were obtained by model checkers that use multiple optimization techniques and, in most cases, also compact symbolic representation of states and transitions in the model.

We also note that $\mathbf{AE}\mu\mathbf{C}$ model checking was distinctly faster than that of $\mathbf{ATL}_{\mathrm{ir}}$ even when they were both (theoretically) NP-complete, i.e., for coalitions larger than 2 agents (cf. Figure 2).

4 Fixpoint Approximation of Strategic Ability: Semantics

In this section, we empirically compare the output of model checking $\mathbf{ATL}_{\mathrm{ir}}$ with that of model checking $\mathbf{AE}\mu\mathbf{C}$. That is, we compare the truth values of $\mathbf{ATL}_{\mathrm{ir}}$ specifications in the benchmark models, versus the truth values of their naive fixpoint approximations. The output of $\mathbf{AE}\mu\mathbf{C}$ model checking has been produced by our model checking algorithm. The truth values according to the $\mathbf{ATL}_{\mathrm{ir}}$ semantics were determined by hand.

Configuration	$\mathbf{AE}\mu\mathbf{C}$ (ψ_2')	$\mathbf{ATL}_{\mathrm{ir}}$ (ψ_2)
4 (1,1,1)	false	false
5 (1,1,2)	false	false
6 (2,1,2)	false	false
6 (3,1,1)	false	false
7 (2,2,2)	false	false
8 (3,2,2)	false	false

Fig. 7. Model checking output for Castles, formula ψ_2' vs. ψ_2

Horses	$\mathbf{AE}\mu\mathbf{C}$ (ϕ_1')	$\mathbf{ATL}_{\mathrm{ir}}$ (ϕ_1)
3	false	false
4	false	false
5	false	false
6	false	false
7	false	false
8	false	false

Fig. 8. Model checking output for TianJi, formula ϕ_1' vs. ϕ_1

4.1 Benchmarks 1 & 2: Castles and TianJi

The results for Castles are presented in Figures 6 and 7. The tables present the truth values of a given formula in the initial state of the benchmark model of a given size. Similarly, the output of model checking for TianJi is presented in Figures 8 and 9.

4.2 Benchmark 3: Modified TianJi

The experiments with Castles have shown that the $\mathbf{AE}\mu\mathbf{C}$ approximations capture a much more restrictive notion of ability than the original $\mathbf{ATL}_{\mathrm{ir}}$ specifications. In this view, the results in Figures 8 and 9 are not very informative. Both semantics have produced the same truth values, but was it because they are indeed so close? Or rather because the $\mathbf{ATL}_{\mathrm{ir}}$ semantics incidentally produced "false," i.e., the truth value that the fixpoint semantics seems to favor?

To answer this question, we have prepared and executed an additional run of experiments based on a modification of the TianJi story. In "Modified TianJi," general TianJi always sees the horse selected by the king before sending his own horse to the next race. The modification significantly increases the strategic abilities of the general. The results are presented in Figure 10.

4.3 Discussion

The experiments show that using $\mathbf{AE}\mu\mathbf{C}$ to approximate model checking of $\mathbf{ATL}_{\mathrm{ir}}$ in a straightforward way does not work. There is no correlation between the truth of the $\mathbf{ATL}_{\mathrm{ir}}$ formulae that we have tested, and their naive $\mathbf{AE}\mu\mathbf{C}$

Horses	**AEμC** (ϕ_2')	**ATL$_{\mathrm{ir}}$** (ϕ_2)
3	false	false
4	false	false
5	false	false
6	false	false
7	false	false
8	false	false

Fig. 9. Model checking output for TianJi, formula ϕ_2' vs. ϕ_2

Horses	**AEμC** (ϕ_1')	**ATL$_{\mathrm{ir}}$** (ϕ_1)
3	false	true
4	false	true
5	false	true
6	false	true
7	false	true
8	false	true

Horses	**AEμC** (ϕ_2')	**ATL$_{\mathrm{ir}}$** (ϕ_2)
3	false	true
4	false	true
5	false	true
6	false	true
7	false	true
8	false	true

Fig. 10. Model checking output for Modified TianJi

translations. In fact, the latter ones did not hold in an overwhelming majority of models that we looked at. This is due to the fact that *persistent* or *recomputable* strategic ability is a much stronger property than being able to come up with a winning strategy only in the initial state of the game.

This suggests two possible ways of further study. One is to identify subclasses of concurrent game structures where recomputable strategies can be obtained. Models of agents with perfect recall seem a natural candidate in this respect. The other is to suitably weaken the **AEμC** translations so that they capture also existence of (some) non-recomputable strategies. We leave exploration of both paths for future research.

5 Conclusions

In this paper, we have looked at verification of strategic abilities for agents with imperfect information. The aim was to investigate whether straightforward fixpoint approximations provide an interesting alternative to formulae of **ATL$_{\mathrm{ir}}$**, for which model checking is known to be theoretically and practically hard. The answer is both yes and no. On one hand, our experimental results show that verification of "fixpoint abilities," specified in alternating epistemic μ-calculus, offers a dramatic speedup over model checking of **ATL$_{\mathrm{ir}}$**. On the other hand, there is no correlation between satisfaction of the **ATL$_{\mathrm{ir}}$** formulae that we have tested, and their naive **AEμC** translations. Thus, we conclude that **AEμC** *is* an attractive alternative to **ATL$_{\mathrm{ir}}$** from the computational point of view, but it does *not* approximate model checking of **ATL$_{\mathrm{ir}}$** in a straightforward way. Rather, it is underpinned by a distinctly different notion of ability, based on existence of *persistent* or *recomputable* strategies.

In the future, we plan to identify subclasses of concurrent game structures where such recomputable strategies can be obtained. We will also investigate how to weaken the fixpoint translations so that they capture also existence of some non-recomputable strategies.

References

1. T. Ågotnes. A note on syntactic characterization of incomplete information in ATEL. In *Procedings of Workshop on Knowledge and Games*, pages 34–42, 2004.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
3. N. Bulling and W. Jamroga. Alternating epistemic mu-calculus. In *Proceedings of IJCAI-11*, pages 109–114, 2011.
4. N. Bulling and W. Jamroga. Comparing variants of strategic ability: How uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
5. S. Busard, C. Pecheur, H. Qu, and F. Raimondi. Improving the model checking of strategies under partial observability and fairness constraints. In *Formal Methods and Software Engineering*, volume 8829 of *Lecture Notes in Computer Science*, pages 27–42. Springer, 2014.
6. Simon Busard and Charles Pecheur. Pynusmv: Nusmv as a python library. In *Proceedings of NASA Formal Methods*, pages 453–458, 2013.
7. C. Dima and F.L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
8. X. Huang and R. van der Meyden. Symbolic model checking epistemic strategy logic. In *Proceedings of AAAI*, pages 1426–1432, 2014.
9. W. Jamroga. Some remarks on alternating temporal epistemic logic. In B. Dunin-Keplicz and R. Verbrugge, editors, *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003)*, pages 133–140, 2003.
10. W. Jamroga and T. Ågotnes. Modular interpreted systems: A preliminary report. Technical Report IfI-06-15, Clausthal University of Technology, 2006.
11. W. Jamroga and J. Dix. Model checking ATL_{ir} is indeed Δ_2^P-complete. In *Proceedings of EUMAS'06*, volume 223 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2006.
12. W. Jamroga and J. Dix. Model checking abilities of agents: A closer look. *Theory of Computing Systems*, 42(3):366–410, 2008.
13. W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
14. A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: An open-source model checker for the verification of multi-agent systems. *International Journal on Software Tools for Technology Transfer*, 2015. To appear.
15. J. Pilecki, M.A. Bednarczyk, and W. Jamroga. Synthesis and verification of uniform strategies for multi-agent systems. In *Proceedings of CLIMA XV*, volume 8624 of *Lecture Notes in Computer Science*, pages 166–182. Springer, 2014.
16. J. Pilecki, M.A. Bednarczyk, and W. Jamroga. Synthesis and verification of uniform strategies for multi-agent systems. 2016. Journal version, under submission.
17. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.